MNRAS **000**, 000–000 (0000) Preprint 15 November 2016 Compiled using MNRAS LATEX style file v3.0

Brute-Force Mapmaking with Compact Interferometers: A MITEoR Northern Sky Map from 128 MHz to 175 MHz

H. Zheng[1⋆], M. Tegmark[1], J. S. Dillon[1,2], A. Liu[2†], A. R. Neben[1], S. M. Tribiano[1,3,4] R. F. Bradley[5,6], V. Buza[1,7], A. Ewall-Wice[1], H. Gharibyan[1,8], J. Hickish[9], E. Kunz[1], J. Losh[1], A. Lutomirski[1], E. Morgan[1], S. Narayanan[1], A. Perko[1,8], D. Rosner[1], N. Sanchez[1], K. Schutz[1,10], M. Valdez[1,11], J. Villasenor[1], H. Yang[1,8], K. Zarb Adami[3], I. Zelko[1,7], K. Zheng[1]

[1] *Dept. of Physics and MIT Kavli Institute, Massachusetts Institute of Technology, 77 Massachusetts Ave., Cambridge, MA 02139, USA*
[2] *Dept. of Astronomy and Radio Astronomy Lab, University of California, Berkeley, CA 94720, USA*
[3] *Science Dept. Borough of Manhattan Community College, City University of New York, New York, NY 10007, USA*
[4] *Dept. of Astrophysics, American Museum of Natural History, New York, NY 10024*
[5] *Dept. of Electrical and Computer Engineering, University of Virginia, Charlottesville, VA 22904, USA*
[6] *National Radio Astronomy Observatory, Charlottesville, VA 22903, USA*
[7] *Dept. of Physics, Harvard University, Cambridge, MA 02138, USA*
[8] *Dept. of Physics, Stanford University, Stanford, CA 94305, USA*
[9] *Dept. of Physics, University of Oxford, Oxford, OX1 3RH, United Kingdom*
[10] *Berkeley Center for Theoretical Physics, University of California, Berkeley, CA 94720, USA*
[11] *Dept. of Astronomy, Boston University, Boston, MA 02215, USA*

arXiv:1605.03980v3 [astro-ph.IM] 12 Nov 2016

15 November 2016

ABSTRACT

We present a new method for interferometric imaging that is ideal for the large fields of view and compact arrays common in 21 cm cosmology. We first demonstrate the method with simulations for two very different low frequency interferometers, the Murchison Widefield Array (MWA) and the MIT Epoch of Reionization (MITEoR) Experiment. We then apply the method to the MITEoR data set collected in July 2013 to obtain the first northern sky map from 128 MHz to 175 MHz at ∼2° resolution, and find an overall spectral index of -2.73 ± 0.11. The success of this imaging method bodes well for upcoming compact redundant low-frequency arrays such as HERA. Both the MITEoR interferometric data and the 150 MHz sky map are available at http://space.mit.edu/home/tegmark/omniscope.html.

Key words: Cosmology: Early Universe – Radio Lines: General – Techniques: Interferometric – Methods: Data Analysis

1 INTRODUCTION

Mapping neutral hydrogen throughout our universe via its redshifted 21 cm line offers a unique opportunity to probe the cosmic "dark ages", the formation of the first luminous objects, and the Epoch of Reionization (EoR). In recent years a number of low-frequency radio interferometers designed to probe the EoR have been successfully deployed, such as the Low Frequency Array (LOFAR; Röttgering 2003), the Murchison Widefield Array (MWA; Tingay et al. 2013), the Donald C. Backer Precision Array for Probing the Epoch of Reionization (PAPER; Parsons et al. 2010), the 21 cm Array (21CMA; Wu 2009), and the Giant Metrewave Radio Telescope (GMRT; Paciga et al. 2011). Unfortunately, the cosmological 21 cm signal is so faint that none of the current experiments have detected it yet, although increasingly stringent upper limits have recently been placed (Paciga et al. 2013; Dillon et al. 2014; Parsons et al. 2014; Jacobs et al. 2015; Ali et al. 2015). Fortunately, the next generation Hydrogen Epoch of Reionization Array (HERA; Pober et al. 2014) is already being commissioned and larger future arrays, such as the Square Kilometre Array (SKA; Mellema et al. 2013), are in the planning stages.

Mapping diffuse structure is important for EoR science. A major challenge in the field is that foreground contamination is perhaps four orders of magnitude larger in brightness temperature than the cosmological hydrogen sig-

⋆ E-mail: jeff.h.zheng@gmail.com
† Hubble Fellow

© 0000 The Authors

nal (de Oliveira-Costa et al. 2008; Dillon et al. 2014; Parsons et al. 2014; Ali et al. 2015). Many first generation experiments therefore focus on a foreground-free region of Fourier space, giving up on any sensitivity to the foreground-contaminated regions (Parsons et al. 2012b; Pober et al. 2013; Liu et al. 2014a,b; Dillon et al. 2015b). To access those regions and regain the associated sensitivity, one must accurately model and remove such contamination, a challenge that requires even greater sensitivity as well as more accurate calibration and beam modeling than the current state-of-the-art in radio astronomy (see Furlanetto et al. 2006; Morales & Wyithe 2010 for reviews). Moreover, an accurate sky model is important for calibrating these low-frequency arrays. For non-redundant arrays—arrays with few or no identical baselines—such as the MWA, modeling the diffuse structure is necessary for calibrating short baselines (Pober et al. 2016; Barry et al. 2016). For redundant arrays such as PAPER, MITEoR, and HERA (Pober et al. 2014), although they can apply redundant calibration which solves for per antenna calibration gains without any sky models (Zheng et al. 2014; Ali et al. 2015), they do need a model of the diffuse structure to lock the overall amplitude of their measurements (Jacobs et al. 2013).

The Global Sky Model (GSM; de Oliveira-Costa et al. 2008) is currently the best model available for diffuse Galactic emission at EoR frequencies. It has been widely used in the EoR community as a foreground model. The sky maps used in the GSM that are close to the EoR frequency (100 MHz to 200 MHz) are the 1999 MRAO+JMUAR map (Alvarez et al. 1997; Maeda et al. 1999) at 45 MHz and 3.6° resolution, and the 1982 Haslam map (Haslam et al. 1981, 1982; Remazeilles et al. 2015) at 408 MHz and 0.8° resolution. Other notable sky maps include the 1970 Parkes maps (Landecker & Wielebinski 1970) near the equatorial plane at 85 MHz and 150 MHz, with 3.5° and 2.2° resolution, respectively. However, both the frequency and sky coverage of these maps are insufficient to achieve the precision necessary for 21 cm cosmology.

All the maps listed above were made using steerable single dish antennas. The sensitivity necessary for EoR science requires a large collecting area, for which steerable single dish radio telescopes become prohibitively expensive (Tegmark & Zaldarriaga 2009a). As a result, the aforementioned EoR experiments have all opted for interferometry, combining a large number of independent antenna elements which are (except for GMRT) individually more affordable. Since these low-frequency instruments are designed for EoR science, they differ in many ways from traditional radio interferometers. For example, PAPER's array layout maximizes redundancy for optimal sensitivity and calibration (Parsons et al. 2012a), which makes it unsuitable for standard interferometric imaging. MWA, on the other hand, has a minimally redundant layout suitable for imaging, and customized algorithms have been developed to optimize its imaging performance such as Sullivan et al. (2012) and Offringa et al. (2014). Some progress has been made to create maps of large-scale structure (Wayth et al. 2015), they have relied on traditional algorithms designed for narrow-field imaging.

Fundamentally, interferometric imaging is a linear inversion problem, since the visibility data are linearly related to the sky temperature: the vector of measured data equals the vector of sky temperature times some huge matrix, with added noise. A number of map-making frameworks using this matrix approach have been proposed in recent years. For example, Shaw et al. (2014) proposed a technique on spherical harmonics designed for drift scan instruments. Ghosh et al. (2015) focused mostly on imaging fields of view on the scale of a few degrees. While Dillon et al. (2015a) develops much of the same formalism that we use, they focus less on high-fidelity mapmaking and more on propagating mapmaking effects into 21 cm power spectrum estimation. The "brute force" matrix inversion and full-sky imaging approach we pursue is desirable since it is optimal in data reduction and conceptually simple, but is very computationally expensive. In terms of computations, in order to have 1° pixel size, the entire sky consists of about 5×10^4 pixels and thus requires the inversion of a 5×10^4 by 5×10^4 matrix. While not computationally prohibitive, it is desirable to minimize the number of such inversions. Full-sky inversion has only becomes more relevant recently with the deployment of low-frequency radio interferometers with very short baselines (on the order of a few wavelengths), which are necessary to probe the structures on large angular scales.

In this work, we propose a practical "brute-force" matrix-based method for imaging diffuse structure using interferometers. We pixelize the temperature distribution in the entire sky into a vector $\boldsymbol{s}$, and flatten all of the complex visibility data from all baselines and all times (as much as a whole sidereal day) into another vector $\boldsymbol{v}$. These two vectors are related by a set of linear equations described by a large matrix $\mathbf{A}$, determined by the primary beam and the fringe patterns of various baselines. We then deploy familiar tools for solving linear equations such as regularization and Wiener filtering to obtain a solution for the entire sky. This algorithm can potentially be applied to any interferometer, and it is especially advantageous for imaging diffuse structure rather than compact sources. The algorithm naturally synthesizes data from different times, without requiring any approximations such as a flat sky, a co-planar array, or a "w-term". It can therefore fill in "missing modes" that are otherwise undetermined using each snapshot individually. Furthermore, it allows for the optimal combination of data from different instruments in visibility space, rather than image space. Lastly, even for arrays with very high angular resolution for which the matrix inversion is impossible, this method can be used to calibrate its short baselines without needing an external sky model.

The rest of this paper is structured as follows. In Section 2, we describe our method in detail. In Section 3, we quantify the accuracy of the method using simulations for both MITEoR's maximally redundant layout and MWA's minimally redundant layout to provide intuitive understanding of its performance. In Section 4, we apply the algorithm to our MITEoR data set to perform absolute calibration, and produce a northern sky map with about 2° resolution centered around 150 MHz. We use the obtained map to compute spectral indices throughout our frequency range, as well as spectral indices when compared to 85 MHz and 408 MHz maps. To estimate the accuracy of our map, we compare it with both the 1970 Parkes map at 150 MHz and the prediction of the GSM at 150 MHz.

2 WIDE FIELD INTERFEROMETRIC IMAGING

2.1 Framework

We start with the interferometer equation for the visibility across the ith and jth antennas, following the conventions in Tegmark & Zaldarriaga (2009b):

$$v_{ij} = \int T_s(\boldsymbol{k}) B(\boldsymbol{k}) e^{i\boldsymbol{k}\cdot\boldsymbol{r}_{ij}}\, d\Omega, \tag{1}$$

where T_s is the sky temperature, B is the primary beam strength, $\boldsymbol{r}_{ij}$ is the baseline vector between the ith and jth antennae, and $\boldsymbol{k}$ is the wave vector of the electromagnetic radiation whose amplitude $|\boldsymbol{k}| = \frac{2\pi}{\lambda}$ and whose direction $\hat{\boldsymbol{k}}$ points towards the observer (Thompson et al. 2001). In principle, one can generalize Eq. (1) by including the full Stokes vector for the polarized sky temperature and a combination of Mueller matrices and Jones matrices for the primary beam, and follow the same algorithm described below to capture full polarization information. In this work, however, we limit our treatment to unpolarized sky temperature. We thus assume that the sky is unpolarized, and estimate its intensity using xx and yy polarization visibilities, which have different beam patterns. As such, there is potential polarization leakage from Stokes Q and U into I, but our wide primary beam and long observation time should limit such leakage.

Radio interferometry typically takes advantage of this equation by first performing a coordinate transformation on Eq. (1) from $\boldsymbol{k}$ on the spherical surface to its projection on the xy-plane:

$$v_{ij} = v(\boldsymbol{u}_{ij}) = \int \frac{T_s(\boldsymbol{q}) B(\boldsymbol{q})}{\sqrt{1-|\boldsymbol{q}|^2}} e^{i2\pi\boldsymbol{q}\cdot\boldsymbol{u}_{ij}}\, d\boldsymbol{q}, \tag{2}$$

where $\boldsymbol{q} = \frac{\lambda}{2\pi}(k_x, k_y)$, and $\boldsymbol{u}_{ij} = \frac{\boldsymbol{r}_{ij}}{\lambda}$ (Tegmark & Zaldarriaga 2009a). By treating the visibilities, v_{ij}, and the beam-weighted sky image, $\frac{T_s(\boldsymbol{q})B(\boldsymbol{q})}{\sqrt{1-|\boldsymbol{q}|^2}}$, as a Fourier pair and performing 2D Fourier transforms on measured v_{ij}'s, radio interferometers have obtained high quality images of the sky with tremendous success. However, this Fourier approach comes with one important limitation. Generally speaking, the beam-weighted sky image is band-limited to the unit circle $|\boldsymbol{q}| \leqslant 1$, so by the Nyquist theorem, one has to have uv sampling finer than half a wavelength to avoid aliasing in the image. In reality, it is difficult to have any baselines shorter than half a wavelength due to the physical size of the antennas. What is more, since the shortest baseline has to be longer than the diameter of the antenna, the largest angular scale available is always smaller than the primary beam width, making aliasing inevitable. Mature techniques such as anti-aliasing filters (Taylor et al. 1999) have been developed to solve this problem, but they are typically tailored for resolving compact structures rather than a diffuse background. The traditional uv-plane approach is therefore not ideal for imaging larger angular scales than the primary beam width.

To overcome this challenge, we follow Dillon et al. (2015a) and attack Eq. (1) from a different angle. By discretizing the integral over sky angles into a sum over sky pixels indexed by n, and including measured visibilities from all times (such as 8 hours during a night's drift scan), Eq. (1) becomes

$$v_{ijt} = \sum_n T_s(\boldsymbol{k}_n) B(\boldsymbol{k}_n, t) e^{i\boldsymbol{k}_n\cdot\boldsymbol{r}_{ij}(t)} \Delta\Omega, \tag{3}$$

where $\Delta\Omega$ is the pixel angular size, and we express all quantities in equatorial coordinates. Here the sky, T_s, is static, and $B(\boldsymbol{k}_i, t)$ and $\boldsymbol{r}_{ij}(t)$ change due to Earth's rotation and possibly the instrument's re-pointing such as in MWA's case. We then flatten visibilities from all times and baselines into one vector $\boldsymbol{v}$, flatten $T_s(\boldsymbol{k}_n)$ for all discretized directions in the sky into $\boldsymbol{s}$, and package the rotating beam and baseline information into a big matrix $\mathbf{A}$. Eq. (3) now takes the form of a system of linear equations

$$\boldsymbol{v} = \mathbf{A}\boldsymbol{s} + \boldsymbol{n}, \tag{4}$$

where we have now taken into account visibility noise $\mathbf{n}$ with mean zero and covariance $\langle \boldsymbol{n}\boldsymbol{n}^t \rangle \equiv \mathbf{N}$. Since the sky temperature is always real whereas $\mathbf{A}$ and $\boldsymbol{v}$ are complex, we "realize" the system by appending the imaginary part of $\mathbf{A}$ and $\boldsymbol{v}$ after their real parts, and double the noise variance.

To optimally estimate $\boldsymbol{s}$, we use the minimum variance estimator (Tegmark 1997)

$$\hat{\boldsymbol{s}} = (\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A})^{-1}\mathbf{A}^t\mathbf{N}^{-1}\boldsymbol{v}, \tag{5}$$

for which the error covariance matrix is

$$\mathbf{\Sigma} \equiv \langle (\hat{\boldsymbol{s}} - \boldsymbol{s})(\hat{\boldsymbol{s}} - \boldsymbol{s})^t \rangle = (\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A})^{-1}. \tag{6}$$

While this is an elegant set-up with a straightforward solution, there are two major technical difficulties with this approach: the size of $\mathbf{A}$ and the invertibility of $\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A}$. We address these topics in detail in the next two sections.

2.2 Constructing the A-Matrix

To obtain an intuitive understanding of the linear system described by Eq. (4), we express the nth column of $\mathbf{A}$ according to Eq. (3), and change the coordinate system back to a rotating sky with a static beam and baselines:

$$\begin{aligned} a_{ijt}^{(n)} &= B(\boldsymbol{k}_n, t) e^{i\boldsymbol{k}_n\cdot\boldsymbol{r}_{ij}(t)} \Delta\Omega \\ &= B(\boldsymbol{k}_n) e^{i\boldsymbol{k}_n(t)\cdot\boldsymbol{r}_{ij}} \Delta\Omega. \end{aligned} \tag{7}$$

For a given baseline pair ij, $a_{ijt}^{(n)}$ is the visibility over time on this baseline produced by the nth pixel in the sky if it had unit temperature. Stacking all baseline pairs together, the nth column of $\mathbf{A}$ is simply the set of visibilities the instrument would have measured if the sky only consisted of a single nth pixel of unit flux, with 0 in all other pixels. Because the visibilities we measure are a linear superposition of contributions from all directions in the sky (Dillon et al. 2015a), finding an optimal solution to Eq. (4) is simply asking how much flux is needed in each pixel of the sky in order to jointly produce the visibilities we actually measured.

With this intuition in mind, we turn to the topic of how to pixelize the sky. Since $\mathbf{A}$ is not sparse, computing Eq. (5) will inevitably involve some form of inversion of $\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A}$, whose computational cost scales as $n_{\rm pix}^3$. We would therefore like the number of pixels to be as small as possible. On the other hand, in order to preserve accuracy in discretizing Eq. (1) to Eq. (3), $\Delta\Omega$ needs to be smaller than square of the angular resolution of the instrument, where the angular resolution is roughly $\frac{\lambda}{d_{\rm max}}$, and $d_{\rm max}$ is the longest baseline

length. When we pixelize the sky using the HEALPIX convention (Górski et al. 2005), we thus need $n_{\rm side} > \frac{d_{\rm max}}{\lambda}$. For MITEoR, whose longest baseline is only about 15 wavelengths, we choose $n_{\rm side} = 64$, which translates to about 4.8×10^4 pixels, which is about the largest size that can be comfortably manipulated on a personal computer. To obtain higher resolution maps for other instruments with even longer baselines, one may choose to use non-uniform pixelization, which we discuss in Appendix B. It is worth pointing out that resolution alone does not decide the number of pixels needed; rather, what matters is the ratio of the primary beam width to the pixel width. Thus, the pixel number may not be very large if one wishes to image a small patch of sky at higher resolution with a narrow-beam instrument, depending on how quickly the sidelobes of the beam fall off away from its center.

2.3 Regularization and Point Spread Functions

For interferometric data sets, $\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A}$ can be poorly conditioned and numerically non-invertible because the instrument is insensitive to certain linear combinations of the sky, so we insert a regularization matrix $\mathbf{R}$ (which we assume to be symmetric throughout this work) into Eq. (5):

$$\hat{\boldsymbol{s}}^{\mathbf{R}} = (\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A} + \mathbf{R})^{-1}\mathbf{A}^t\mathbf{N}^{-1}\boldsymbol{v}. \tag{8}$$

We then substitute Eq. (4) into Eq. (8), which leads to

$$\begin{aligned}\hat{\boldsymbol{s}}^{\mathbf{R}} =&(\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A} + \mathbf{R})^{-1}\mathbf{A}^t\mathbf{N}^{-1}(\mathbf{A}\boldsymbol{s} + \boldsymbol{n})\\ =&\mathbf{P}\boldsymbol{s} + (\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A} + \mathbf{R})^{-1}\mathbf{A}^t\mathbf{N}^{-1}\boldsymbol{n}, \qquad (9)\\ \langle\hat{\boldsymbol{s}}^{\mathbf{R}}\rangle =&\mathbf{P}\boldsymbol{s} \qquad (10)\end{aligned}$$

where we define the point spread matrix $\mathbf{P} \equiv (\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A} + \mathbf{R})^{-1}\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A}$. The matrix $\mathbf{P}$ is not the identity matrix due to the insertion of $\mathbf{R}$, so each column of $\mathbf{P}$ acts as a point spread function (PSF) for the corresponding pixel in the true sky map $\boldsymbol{s}$. The introduction of regularization thus represents a comprise between the goal of completely removing the PSF and producing a map with favorable noise properties. The result is, in the language of traditional radio astronomy, a position-dependent synthesized beam. For each of the PSF, we can calculate the effective full width half maximum (FWHM) using

$$\theta_{\rm FWHM} = \sqrt{\pi^{-1} n_{\frac{1}{2}} \Delta\Omega}, \tag{11}$$

where $n_{\frac{1}{2}}$ is the number of pixels in the PSF whose absolute value is above half of the maximum pixel, and $\Omega_{n_{\rm side}}$ is the angular area for each pixel. We use $\theta_{\rm FWHM}$ to represent angular resolution throughout this work.

It is also worth noting that the sum of each row of $\mathbf{P}$ is a damping factor for the solution. For a fictitious uniform sky, if each row of $\mathbf{P}$ does not sum to unity, then the resulting solution will not retain the same uniform amplitude as the original fictitious sky. Thus, it is desirable to renormalize each row of $\mathbf{P}$ to sum to unity.

With the introduction of $\mathbf{R}$, the error covariance is now

$$\begin{aligned}\mathbf{\Sigma}^{\mathbf{R}} =&\langle(\hat{\boldsymbol{s}}^{\mathbf{R}} - \mathbf{P}\boldsymbol{s})(\hat{\boldsymbol{s}}^{\mathbf{R}} - \mathbf{P}\boldsymbol{s})^t\rangle\\ =&(\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A} + \mathbf{R})^{-1}(\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A})(\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A} + \mathbf{R})^{-1}\\ =&\mathbf{P}(\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A} + \mathbf{R})^{-1}. \qquad (12)\end{aligned}$$

To quantify how well the predicted error in Eq. (12) accounts for actual noise we obtain in either our simulations or maps using real data, we define the normalized error δ at the ith pixel as

$$\delta_i = \frac{|\hat{s}_i^{\mathbf{R}} - (\mathbf{P}\mathbf{s})_i|}{\sqrt{\Sigma_{ii}}}. \tag{13}$$

Intuitively, δ represents the discrepancy between the recovered map and the ground truth in units of the expected noise level, and is thus expected to center around 1 in absence of any systematic effects.

Now we discuss the choice of $\mathbf{R}$. Inserting $\mathbf{R}$ is equivalent to having prior knowledge of the sky, where the Bayesian prior for $\mathbf{s}$ has mean $\langle\mathbf{s}\rangle = \boldsymbol{s}^p = 0$ and covariance matrix $\langle(\boldsymbol{s} - \boldsymbol{s}^p)(\boldsymbol{s} - \boldsymbol{s}^p)^t\rangle = \mathbf{R}^{-1}$. If one has a sky model with well-characterized error properties (not the case for this work), there is a natural choice of $\mathbf{R}$ which we discuss in the next section. In the absence of such a model, however, the choice of $\mathbf{R}$ depends on the array layout, noise properties of the visibilities, as well as the trade-off between angular resolution and noise (we will demonstrate these qin much more detail through simulation in Section 3). The simplest form of $\mathbf{R}$ is a uniform regularization matrix proportional to the identity: $\mathbf{R} = \epsilon^2\mathbf{I}$. ϵ^{-1} has the same units as $\boldsymbol{s}$, and can be compared to the noise level in the map solution: smaller ϵ^{-1} suppresses noisy modes more strongly since it will dominate the diagonal of $\mathbf{\Sigma}^{\mathbf{R}}$, but it also hurts angular resolution as it introduces wider point spread functions by bringing $\mathbf{P}$ farther away from the identity matrix. Therefore, ϵ is a tunable parameter deciding the trade-off between resolution and noise. In our simulations in Section 3 and map making using MITEoR data in Section 4, we show various choices of ϵ, and we leave investigations of the optimal choice of $\mathbf{R}$ for future work.

2.4 Wiener Filtering and Incorporating Prior Knowledge

If we have prior knowledge of the sky such as previous measurements, we would like to optimally combine the existing sky map $\boldsymbol{s}^p$, with our visibility measurements (this is not carried in our simulations or our final map product). We accomplish this by shifting our focus from $\boldsymbol{s}$ to $\boldsymbol{s} - \boldsymbol{s}^p$, and Eq. (5) becomes

$$\hat{\boldsymbol{s}} - \boldsymbol{s}^p = (\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A})^{-1}\mathbf{A}^t\mathbf{N}^{-1}(\boldsymbol{v} - \mathbf{A}\boldsymbol{s}^p). \tag{14}$$

We then quantify the uncertainty in our prior knowledge, by defining the covariance matrix $\mathbf{S} \equiv \langle(\boldsymbol{s} - \boldsymbol{s}^p)(\boldsymbol{s} - \boldsymbol{s}^p)^t\rangle$, and we use $\mathbf{R} = \mathbf{S}^{-1}$ as the regularization matrix:

$$\hat{\boldsymbol{s}}^{\mathbf{R}} - \boldsymbol{s}^p = (\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A} + \mathbf{S}^{-1})^{-1}\mathbf{A}^t\mathbf{N}^{-1}(\boldsymbol{v} - \mathbf{A}\boldsymbol{s}^p). \tag{15}$$

Our regularized estimate for $\boldsymbol{s}$ then becomes

$$\hat{\boldsymbol{s}}^{\mathbf{R}} = (\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A} + \mathbf{S}^{-1})^{-1}\mathbf{A}^t\mathbf{N}^{-1}(\boldsymbol{v} - \mathbf{A}\boldsymbol{s}^p) + \boldsymbol{s}^p. \tag{16}$$

There are three ways of understanding our choice of regularization $\mathbf{R} = \mathbf{S}^{-1}$. First, using the identity

$$(\mathbf{X}^{-1} + \mathbf{Y}^{-1})^{-1} = \mathbf{X}(\mathbf{X} + \mathbf{Y})^{-1}\mathbf{Y} \tag{17}$$

for invertible[1] square matrices $\mathbf{X}$ and $\mathbf{Y}$, one can show that

[1] We have mentioned that $\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A}$ is usually numerically un-

Eq. (15) is equivalent to applying a Wiener filter $\mathbf{W}$ to Eq. (14):

$$\hat{\boldsymbol{s}}^{\mathbf{R}} - \boldsymbol{s}^{p} = \mathbf{W}(\hat{\boldsymbol{s}} - \boldsymbol{s}^{p}), \tag{18}$$

where $\mathbf{W} = \mathbf{S}(\mathbf{\Sigma} + \mathbf{S})^{-1}$, with $\mathbf{\Sigma}$ defined in Eq. (6).

Secondly, by using Eq. (17), one can show that Eq. (16) is equivalent to an inverse variance weighted average of the unregularized $\hat{\boldsymbol{s}}$ and the prior knowledge $\boldsymbol{s}^{p}$:

$$\hat{\boldsymbol{s}}^{\mathbf{R}} = \mathbf{S}(\mathbf{\Sigma} + \mathbf{S})^{-1}\hat{\boldsymbol{s}} + \mathbf{\Sigma}(\mathbf{\Sigma} + \mathbf{S})^{-1}\boldsymbol{s}^{p}, \tag{19}$$

where $\mathbf{\Sigma}$ and $\mathbf{S}$ are the noise covariance matrices for the unregularized $\hat{\boldsymbol{s}}$ and the prior knowledge $\boldsymbol{s}^{p}$, correspondingly. It is reassuring to see from Eq. (19) that $\hat{\boldsymbol{s}}^{\mathbf{R}} \to \boldsymbol{s}^{p}$ when $\mathbf{\Sigma} \gg \mathbf{S}$, and vice versa.

Lastly, Eq. (16) is equivalent to combining visibility data and a previous sky map through appending $\boldsymbol{s}^{p}$ to $\boldsymbol{v}$ and the identity matrix $\mathbf{I}$ to $\mathbf{A}$ in Eq. (4), and solving it using Eq. (5) without any regularization.

Incorporating prior knowledge, especially in the form of existing pixel maps, can indeed fill in the "missing modes" and complement visibility data sets very well. However, we decided to leave demonstration of this part to future work for two reasons. First, there is no existing pixel map at the MITEoR frequency that covers the entire sky region observed by MITEoR. Second, our goal in this work is to demonstrate the effectiveness of our algorithm without access to additional data, despite the "missing mode" problem, so we therefore refrain from resorting to pixel maps to solve the invertibility challenges. Since we have made the data public, we very much hope that other authors will perform such improved regularization in the future.

2.5 Further Generalization

Throughout this section, we have, for clarity, limited our discussion to data sets in the form of baseline by time at a given frequency from one instrument. We generalize this to incorporate data sets from multiple frequencies and even multiple instruments. To synthesize multiple instruments at the same frequency, we simply append all the flattened data vectors together, and stack their corresponding $\mathbf{A}$ matrices in the same order. We then solve for the sky in one step using Eq. (5) or Eq. (8). We demonstrate this through simulation in Section 3.3.

To synthesize multiple frequencies, we assume that the sky map only differs by an overall scaling factor throughout a given frequency range, and the accuracy approximation depends on how wide the frequency range is. If at each frequency ν we have

$$\boldsymbol{v}_{\nu} = \mathbf{A}_{\nu}\boldsymbol{s}_{\nu} \tag{20}$$

together with

$$\boldsymbol{s}_{\nu} = f(\nu)\boldsymbol{s}, \tag{21}$$

we obtain

$$\boldsymbol{v}_{\nu} = f(\nu)\mathbf{A}_{\nu}\boldsymbol{s} = \mathbf{A}'_{\nu}\boldsymbol{s}. \tag{22}$$

invertible, but as long as one excludes pixels never above the horizon, it should be formally invertible, with some eigenvalues very close to zero.

If we know $f(\nu)$, we can simply stack $\boldsymbol{v}_{\nu}$'s and $\mathbf{A}'_{\nu}$'s and solve for $\boldsymbol{s}$. In reality, although we do not know $f(\nu)$ very accurately, we can iterate this, where we start with an estimated $f_0(\nu)$ and at each iteration reevaluate $f(\nu)$ given by the best-fit ratio between $\boldsymbol{v}_{\nu}$ and $\mathbf{A}_{\nu}\boldsymbol{s}$ from the previous iteration. As a result, we can empirically obtain $f(\nu)$ in addition to $\boldsymbol{s}$. We demonstrate this in detail in Section 4.

2.6 Computational Discussion

As mentioned above in Sections 2.2 and 2.3, carrying out the matrix operations in Eq. (8) is computationally intensive. We therefore provide some computational details here for those readers interested in carrying out such calculations.

First of all, it is important to judiciously choose the order of computations when following Eq. (8), since a matrix-vector multiplication is much faster than a matrix-matrix multiplication. For example, it is much faster to calculate $\mathbf{A}^{t}(\mathbf{N}^{-1}\boldsymbol{v})$ than $(\mathbf{A}^{t}\mathbf{N}^{-1})\boldsymbol{v}$.

Matrix multiplication is the dominant computational cost in the entire algorithm, whose cost scales as $n_{\rm pix}^{2}n_{\rm data}$, where $n_{\rm data}$ is the number of data points, or the number of rows in the $\mathbf{A}$-matrix. Since typically $n_{\rm data} > n_{\rm pix}$, this is more expensive than inverting $(\mathbf{A}^{t}\mathbf{N}^{-1}\mathbf{A}+\mathbf{R})$, which scales as $n_{\rm pix}^{3}$. Due to the fact that the $\mathbf{A}$-matrix is neither sparse nor symmetric, iterative methods for linear system of equations, such as conjugate gradient, do not provide significant computational benefit. As a result, direct matrix multiplication of $\mathbf{A}^{t}\mathbf{N}^{-1}\mathbf{A}$ is preferable. Fortunately, matrix multiplication is trivially parallelizable, so parallelizing the matrix multiplication on multiple machines can significantly reduce computation time. The entire computation takes only of order a day on a modern linux computer with sufficient memory to hold the relevant matrices.

In addition, in order to compute the error covariance matrix defined by Eq. (12), direct inversion of $(\mathbf{A}^{t}\mathbf{N}^{-1}\mathbf{A}+\mathbf{R})$ cannot be avoided. In our simulations discussed in the next section, we find that $(\mathbf{A}^{t}\mathbf{N}^{-1}\mathbf{A}+\mathbf{R})$ has non-negligible off-diagonal terms, so approximating the matrix inversion using the inverse of the diagonal is not desirable. However, in circumstances when one does not need the error covariance matrix, such as when experimenting with the choice of $\mathbf{R}$-matrix, one can significantly reduce computational cost by taking advantage of the Cholesky decomposition, which provides a factor of 5-10 speed-up depending on the platform.

Lastly, despite the goal of reducing computational cost, it is important to use double precision for matrix operations when implementing our method. Although neither the data nor the beam pattern are understood to any level near the numerical precision of single point floating numbers, and our simulation results do not rely on such precise knowledge (we calculate $\mathbf{A}$ using single precision, and the simulated data have noise), we found that using single precision during matrix operations has a significant and negative impact on the quality of our results because of the wide dynamic range of eigenvalues.

3 SIMULATIONS

In this section, we perform simulations for both MITEoR and the MWA to demonstrate the algorithms we have described in the previous section. The distinct array layouts of MWA and MITEoR complement each other in demonstrating various aspects of our algorithm. The simulations are all based on just one night of observation on each instrument. To simulate visibilities, we pixelize the GSM using to the HEALPIX resolution $n_{\rm side} = 128$, but when we solve for the sky we use $n_{\rm side} = 32$ (a pixel size of 2°), so that we can quantify any potential errors introduced by having coarse pixels. The noise is simulated as Gaussian noise independent across baselines, frequencies and times, whose amplitude is set to the simulated autocorrelation over $\sqrt{\Delta\nu\Delta t}$, where $\Delta\nu$ is the bandwidth for each frequency bin, and Δt is the integration time. For MITEoR, the noise variance is also scaled down by the redundancy factor (the number of antenna pairs sharing a baseline type) for each baseline type. We then follow Sections 2.1-2.3 to compute maps and their error properties.

3.1 MITEoR Simulation

MITEoR is a highly redundant array. It consists of 64 dual-polarization antennas on a square grid with 3 m spacing. For each polarization, it has 2,016 baselines with 112 unique baseline types. For this simulation, we include the shortest 102 out of its 112 unique baseline types, with baseline lengths between 3 m and 25.5 m. The primary beam model is numerically simulated as described in Zheng et al. (2014), and its FWHM is about 40° near 150 MHz. We simulate for a single night's observation in the local sidereal time (LST) range between 12 hours and 24 hours to resemble the LST coverage in Section 4, at 150 MHz, with 144 seconds integration, and 0.75 MHz frequency bin width. The $\mathbf{A}$-matrix size for this set-up is $122,400 \times 9,725$. We choose $\epsilon^2\mathbf{I}$ as the regularization matrix, where $\epsilon^{-1} = 100$ K (The result is not sensitive to the choice of ϵ within an order of magnitude, as shown later in Fig. 2.).

The results we obtain are shown in the first column of Fig. 1, in which the median noise is 14.9 K. We computed the position-dependent FWHM using the PSF matrix, and found that 96% of the pixels have FWHM less than the pixel size, which means that the resolution is limited by pixelization rather than the PSF. Note that Cygnus A (Cyg A) and Cassiopeia A (Cas A) cast noticeable "ringing" due to their extreme brightness and the effect of the PSF. Due to the coarseness of the pixelization in this section, we defer a demonstration of point source removal, as applied to Cyg A and Cas A, to Section 4. Lastly, the normalized residual map shows that the errors in the recovered map are well described by the noise properties described by the noise covariance matrix $\mathbf{\Sigma^R}$, with the exception of those regions with sharp features, such as pixels near Cyg A, Cas A, and the Galactic center. This is caused by the coarse pixelization, and can be remedied by decreasing the pixel size.

3.2 MWA Simulation

MWA is a minimally redundant array, with 8,128 different baselines. For this simulation, we use the shortest 195 cross-correlation baselines, with baseline lengths between 7.7 m and 25.5 m (referred to as MWAcore from hereon). We limit the baseline lengths in order to use the same $n_{\rm side} = 32$ pixelization as the MITEoR case, as longer baselines requires a much finer pixelization. The primary beam model is obtained by calculating the analytic expression for phased array of short dipoles, and the FWHM is about 14° around 150 MHz. We simulate for a single night's observation in the LST range between 12 hours and 24 hours, at 150 MHz, with 144 seconds integration, and 0.64 MHz frequency bin width. The beam remains in the zenith-pointing drift scan mode for the entire observation, which is not how the MWA typically operates (but doable). The $\mathbf{A}$-matrix size for this set-up is $234,000 \times 9,785$. We choose $\epsilon^2\mathbf{I}$ as the regularization matrix, where $\epsilon^{-1} = 300$ K. The results we obtain are shown in the second column of Fig. 1. The median noise in the map is 47.8 K, and the higher noise lines follow the trajectory of the nulls between the main lobe and the first side lobe, as well as between the first and second side lobes. In terms of the angular resolution, 89% of the sky is limited by pixelization.

3.3 Simulation Discussion

By combining the MWAcore visibilities and the MITEoR visibilities, and using a uniform regularization $\epsilon^{-1} = 100$ K, we obtain another map shown in the third column of Fig. 1, whose median noise is 14.6 K. In this case, 97% of the pixels' resolution is limited by pixelization. To demonstrate the effect of tuning the parameter ϵ in the regularization matrix, we show a few different choices of ϵ in Fig. 2. As the regularization varies from two orders of magnitude too weak to two orders of magnitude too strong, overall noise level decreases, but stronger PSFs start to overly suppress the less sensitive regions near the nulls of the MWA beam and make the temperature negative, and create worse "ringing" around Cyg A and Cas A.

As shown in Fig. 1, with just one night's data, our algorithm can produce high quality diffuse structure maps using either the MWAcore or MITEoR. With the MWAcore's set-up where the shortest baseline is longer than 7 wavelengths, the algorithm successfully determines large scale modes such as the overall amplitude of the map. On the other end, the longest baselines we include are 13 wavelengths long, which naively translates to about 5° resolution, but the matrix approach recovers angular scales smaller than 2°, as both simulations show that the resolution is limited by the 2° pixel size. As we will show in Section 4 using 1° pixels, the true resolving power of these baselines are between 1° and 2°.

Comparing the noise maps of the MWAcore and MITEoR, where MITEoR's 14.9 K median noise level is slightly lower than MWAcore's 47.8 K, we see that the overall noise level is not very sensitive to the number of baseline types, the baselines' length distribution, or the primary beam shape. Although the collecting area of MWA tiles is 16 times larger than each bow-tie antenna used in MITEoR, it is offset by the fact that there are effectively about 2000 baselines used in the MITEoR simulation (with 112 unique baseline types) compared to the MWAcore's 195, so the noise levels are comparable in both cases. In contrast to the overall noise level, the spatial patterns of noise do depend heavily on the

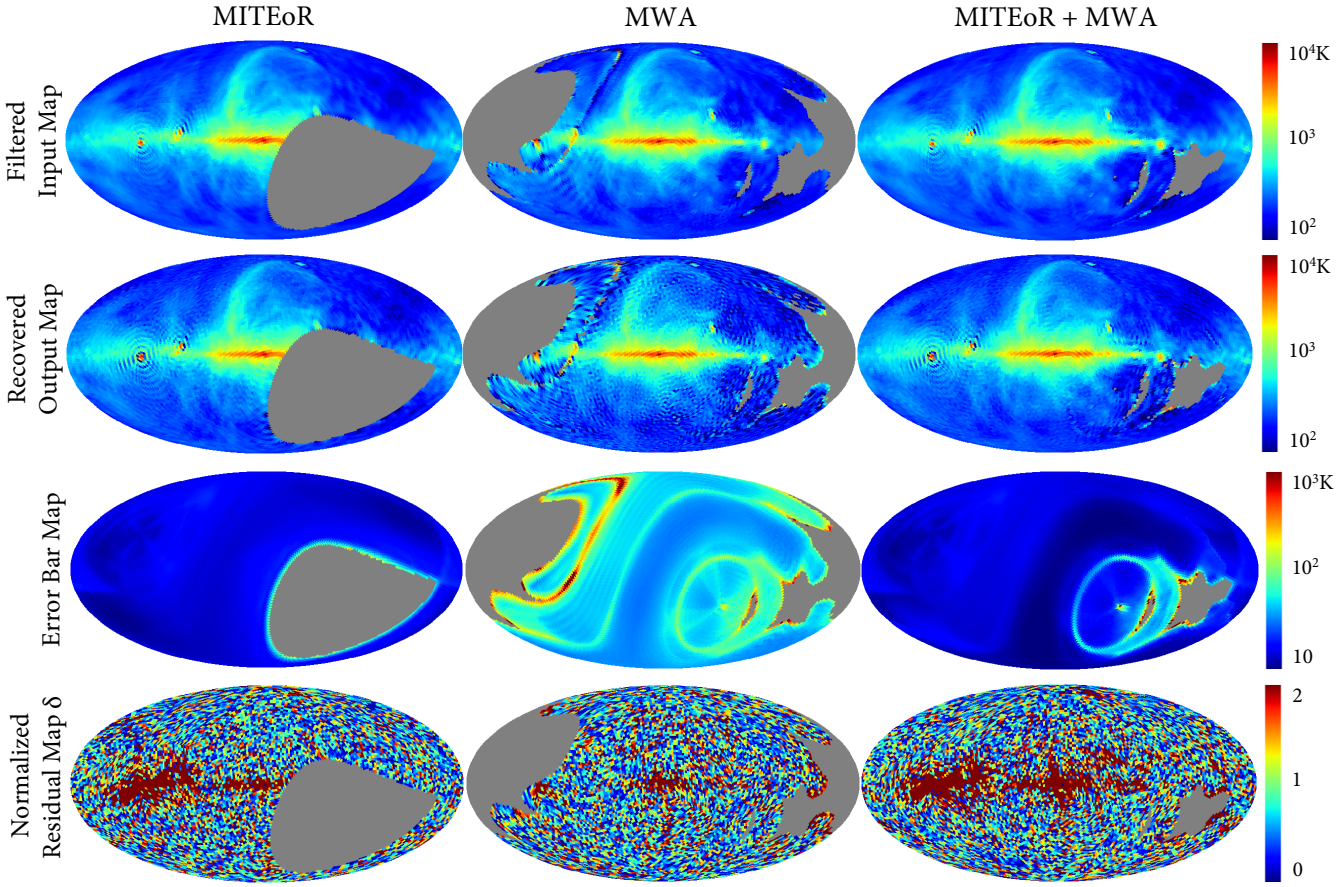

Figure 1. Simulated results for MITEoR (left), the MWAcore (middle), and the two combined (right). Filtered input GSM maps (top) are generated by applying the PSF matrix $\mathbf{P}$ to the GSM with $n_{\rm side} = 32$. Grey areas represent directions that are either never above the horizon or have negligible sensitivity. Recovered output maps are obtained by applying Eq. (8) to simulated noisy visibilities (assuming one night's observation and one frequency channel with < 1 MHz bandwidth), using uniform diagonal regularization matrices with ϵ^{-1} of 100 K, 300 K, and 100 K respectively. The error bar maps are obtained by plotting the square roots of the diagonal entries in $\mathbf{\Sigma^R}$, and the color scale is one order of magnitude smaller than the output maps. Lastly, the normalized residual δ maps represent the ratio of the actual error in our maps to the error bars, as defined in Eq. (13), and their values center around 1 as we expect. Here and in the subsequent figures, all maps are in Mollweide projection centered on the Galactic center.

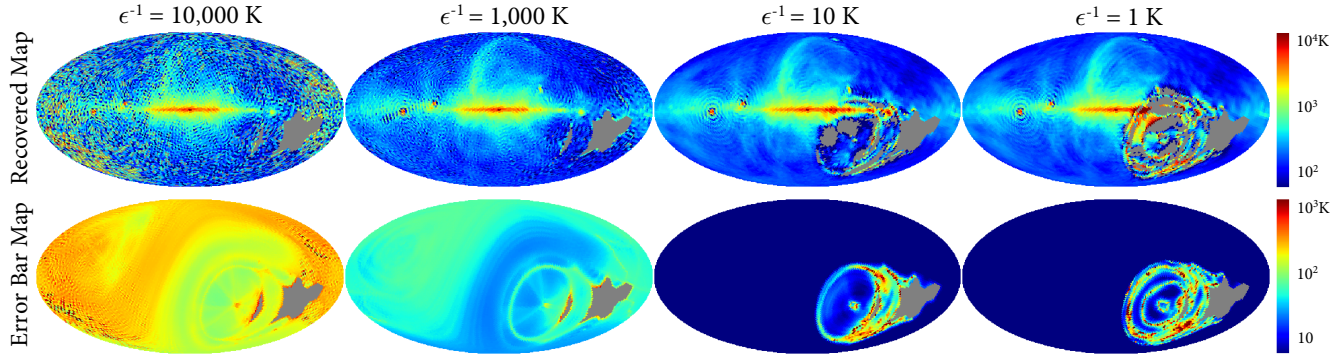

Figure 2. Output maps recovered in simulation and their error bars ($\sqrt{\mathbf{\Sigma^R}_{ii}}$), using the same set of data but different regularization matrices. As ϵ increases over 5 orders of magnitude ($\epsilon^{-1} = 100$ K is shown in the third column of Fig. 1), the noise is significantly suppressed, but the properties of PSF matrix become less desirable: we start to see negative temperatures near the null regions of the MWA beam. This is because stronger regularization suppresses the "missing modes" more heavily, and as the large scale structures near the MWA's side lobe region gets suppressed, more of that region becomes negative. In addition, strong regularization creates worse "ringing" near Cyg A and Cas A, and even near the north pole spur in the 1 K case.

primary beam pattern, which is not surprising, since the primary beam pattern heavily influences the instrument's sensitivity to different parts of the sky.

The δ maps show that the error properties are well characterized by $\mathbf{\Sigma^R}$, with the majority of the δ values less 2. We see that although the visibilities are simulated using the GSM with $n_{\rm side} = 128$, the crude pixelization of $n_{\rm side} = 32$ is not introducing significant errors, other than in regions near Cyg A, Cas A, or the Galactic center. The importance of pixelization errors depends on the baseline lengths included, as well as the amount of noise in the visibilities. We find that pixelization errors become significant if we increase the baseline length threshold by another 50% to about 40 m. We also expect the pixelization error to become dominant if the visibilities have much lower noise, such as when they are averaged over 100 nights as opposed to a single night used here.

Given the result that just one night's data from MWAcore's zenith pointing scan can determine more than half of the sky to better than 50 K precision, MWA has the potential to make a high quality (< 10 K noise through multiple nights' observations) southern sky map with multiple beam pointings and multiple nights throughout the year to fill the sky. In addition, since MWAcore only includes a small fraction of MWA's baselines, a much higher angular resolution (less than a degree) is also achievable with finer pixelizations.

4 NEW SKY MAP

As we have demonstrated in the previous section, the new imaging method works well on simulated data from both MITEoR and MWAcore. In this section, we apply the method to real data collected by MITEoR, and produce a Northern sky map at 150 MHz.

4.1 MITEoR Instrument and Data Reduction

As mentioned above, MITEoR (Zheng et al. 2014) is a compact radio interferometer with 64 dual-pol antennas, deployed in July 2013 in The Forks, Maine (latitude 45.3°). The antennas are identical to individual MWA bow-tie antennas (without beam forming as tiles), and the full width half maximum (FWHM) of the primary beam is about 40° throughout our frequency range. The data used in this work are collected with an 8 by 8 square grid array configuration with 3 m spacings. The 8 bit correlator cross-correlates all 128 antenna-polarizations (each bow-tie antenna has two polarizations as outputs), with integration time of 2.7 seconds, instantaneous bandwidth of 12.5 MHz (tunable between 125 and 185 MHz), and frequency bin width of 50 kHz. The data used in this work were collected through 7 observing sessions, as shown in the left panel of Fig. 3.

We first perform redundant calibration on the raw data, using our redundant calibration pipeline described in Zheng et al. (2014), with further improvements described in Appendix A. Redundant calibration compresses the data in the baseline direction from 2016 cross-correlation visibilities per snapshot to 112 unique baseline types, and automatically flags bad antennas, baselines, frequencies, and time stamps from the data. It is worth noting that redundant calibration uses only the self-consistency between redundant baselines, and does not use any sky model. After redundant calibration, we further compress the data to 0.75 MHz frequency bin width by averaging every 15 frequency bins. We then average over the time direction in 2 minute intervals. We empirically estimate noise during the time averaging step by performing linear fitting over 2 minutes of data and calculating the residual power.

At this stage, we have a data cube of 4 polarizations by 75 frequency bins by 240 time steps by 112 baselines. Since we have not used any sky information, the data is not yet absolute-calibrated, meaning that for each of the $4 \times 75 = 300$ different polarization-frequency data sets, we have 3 undetermined numbers: one overall amplitude, and two re-phasing degrees of freedom. These numbers cannot be determined without performing absolute calibration (as opposed to redundant calibration) with a sky model, which we describe in the next two sections.

4.1.1 Absolute Amplitude Calibration

The goal of absolute calibration is to determine three numbers (an overall amplitude and two re-phasing degrees of freedom) for every time by baseline data set, where each data set consists of more than 10^4 visibility measurements. This is a drastically overdetermined system given a sky model. In this section, we discuss how we determine the overall amplitude. Since we intend to use the map obtained in this work to improve the GSM in a future work, we choose not to use the GSM as our sky model. Rather, we use Cyg A and Cas A as our calibrators. After we obtain a map in Section 4.2, we will lock the amplitude of our map to the Parkes map at 150 MHZ (Landecker & Wielebinski 1970), so the amplitude calibration and its error will only affect our spectral index results, not the amplitude of the map.

Our amplitude calibration is based on extrapolating the frequency-dependent secular decreasing flux models of Cas A (Vinyaikin 2014) and the spectrum of Cyg A from Vinyaĭkin (2006). Cyg A is an ultra-luminous, jet-powered, radio-loud galaxy. For our Cyg A calibration spectrum, we use the model in Eq. 6a in Vinyaĭkin (2006) of a transparent source, with a power law spectrum with spectral index α, observed through the absorbing ionized gas of our Galaxy. The data point used in the model at 152 MHz has a reported 3% error in Parker (1968), and propagating the model parameters' error bars leads to a maximum of 3.2% error in our frequency range.

The frequency dependence of the decreasing flux of Cas A has been widely studied (see Helmboldt & Kassim (2009) and references, therein), and we adopt the empirical model in Vinyaikin (2014) that fits to the accumulated published data taken from 1961 to 2011 from about 12 MHz - 93 GHz, including their most recent observations. The spectrum of Cas A is evaluated in frequency range 125-175 MHz, using the fitted function in Eqs. 15 and 16 in Vinyaikin (2014). With this spectrum modeled at epoch 2015.5, and the model for the frequency dependent secular decrease in Eqs. 9 and 10 of Vinyaikin (2014), we evaluate the spectrum for Cas A during MITEoR observations in August 2013, approximately two years earlier. The largest source of error in the radio spectrum model of Cas A comes from our lack of complete understanding of the behavior of this supernova rem-

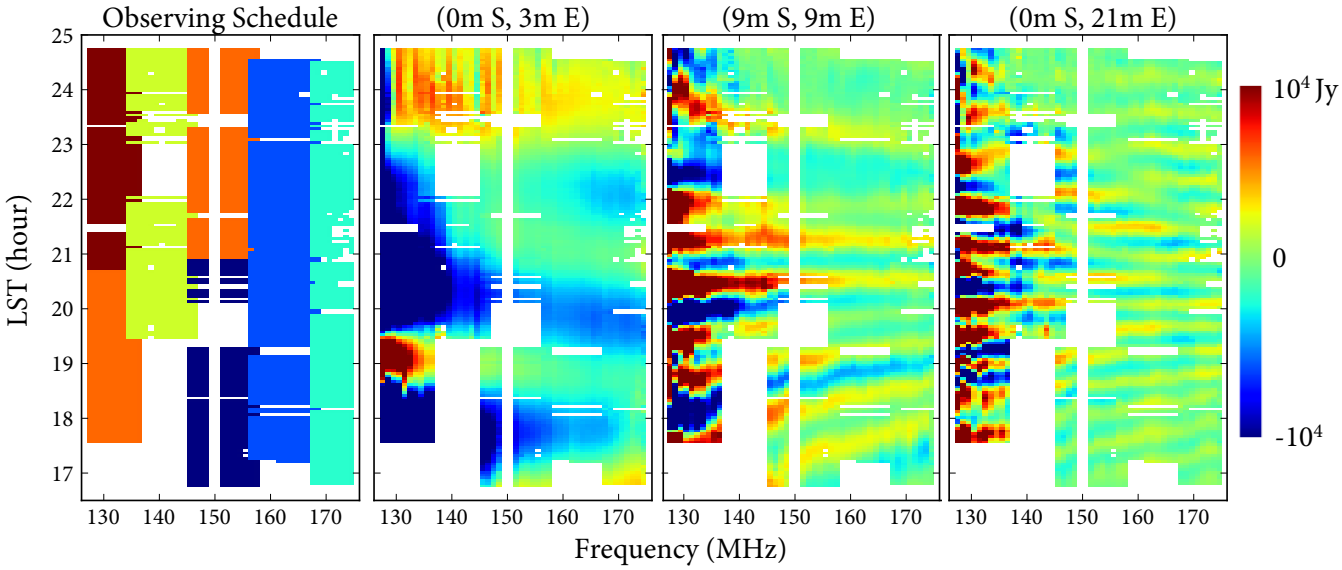

Figure 3. MITEoR's observing schedule on the left and three sets of visibilities in the MITEoR data product on the right. Each of the six colors represent one of the nights between July 27th and August 2nd, 2013. Midnight corresponds to LST at roughly 21 hours. The real part of visibilities over time and frequency on three baselines of very different lengths are shown here. The white gaps in frequency and time are RFI events automatically flagged by the redundant calibration pipeline.

nant. In an evaluation of possible periodic deviations, Helmboldt & Kassim (2009) identify 4 possible modes ranging from 3 to 24 years, in a slightly lower frequency range of interest, $38-80$ MHz, contributing to flux deviations from a secular decrease of up to 15%.

We calibrate in the LST range between 19 and 23 hours, during which Cyg A's elevation ranges from 59° to 85°, and Cas A from 36° to 65°. To minimize error introduced by diffuse structures, we only include for calibration baselines longer than 8.6 wavelengths, which are the longest 9 baselines at the lowest frequency and 34 at the highest frequency. After fitting using Cas A and Cyg A, we find about 15% residual on the visibilities. Since the errors introduced by the Galactic plane are averaged down over different baseline types, we estimate our amplitude calibration to go down by a factor equal to the square root of the number of baselines used. Thus, our absolute calibration has an overall error of about 5% at the lowest frequency and 2.5% at the highest frequency, relative to the calibrators.

4.1.2 Absolute Phase Calibration

As discussed in Zheng et al. (2014), the two re-phasing degrees of freedom (or re-phasing degeneracies) in the visibility space correspond to shifting the beam-weighted sky image $\frac{T_s(\boldsymbol{q})B(\boldsymbol{q})}{\sqrt{1-|\boldsymbol{q}|^2}}$, and cannot be determined using only the self consistency of visibility data without a sky model. However, this is only true for isolated snapshots in time. For instruments with large fields of view such as MITEoR, rotation of the sky does not exactly translate into shifting the beam-weighted sky image in the projected **q**-plane, so a constant shift of the beam-weighted sky image caused by constant re-phasing cannot be consistent with a rotating sky. Therefore, we can use a procedure conceptually similar to self-cal to determine the re-phasing: we first image using the visibility without correcting for the re-phasing degrees of freedom, then use the image we obtain to solve for the re-phasing, and iterate until convergence. In theory, this can be done without any prior sky model, but since each iteration can be computationally expensive, we use the GSM to provide the initial re-phasing solution, and start iterating from there. It is worth noting that, at a given frequency and polarization, unlike self-cal which is solving for, say, 128 antenna calibration parameters at every time stamp, here we are only solving for 2 numbers for an entire observing session, so this iterative algorithm has negligible impact on the validity of Eq. (12). For more detailed discussions of redundant phase calibration, we refer the readers to Zheng et al. (2014) and Appendix A of this work.

In addition, even for instruments whose array layout prevents making usable images, this "self-cal" approach can be applied to remove re-phasing degeneracies. Rather than inserting regularization matrices to make $\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A}$ invertible, if the goal is to calibrate out the re-phasing degeneracies, we can simply use a pseudo-inverse for $\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A}$. For example, for a redundant array with no short baselines, a pseudo-inverse will remove large scale structures in the image, but this has no effect when the resulting image is used as a model to simulate visibilities on those long baselines, so "self-cal" should work just as well.

4.1.3 Cross-talk Removal

We define cross-talk as additive offsets on visibilities that are proportional to the amplitude of auto-correlation, with a small but constant proportionality coefficient. In theory, the cross-talk terms can be solved for in an iterative fashion similar to how we determine the re-phasing degeneracies in the previous section. However, due to the level of thermal noise and systematics present in our data set, cross-talk on our

shortest baselines is highly degenerate with having a bright stripe in the trajectory of our local zenith. Thus for this work, we use the GSM to perform cross-talk removal. We use the GSM to simulate all the visibilities we measure, and for each visibility time series, we solve for the best fit using the GSM model visibility and the auto-correlation, and subtract the auto-correlation component from our data. Thus, for each visibility time series over a few hours, we use the GSM to fit and remove one degree of freedom corresponding to cross-talk.

4.2 Northern Sky Map Combining Multiple MITEoR Frequencies

We apply the algorithm described in Section 2 on the MITEoR data to obtain our Northern sky map. We have shown in Section 3 that the MITEoR data can in principle make high quality maps at individual frequency bins, but due to the systematics present in our data, which we discuss more in the next section, we are not able to make high quality maps using each individual frequency alone. Since in our frequency range the diffuse emission is dominantly synchrotron, which follows a smooth power law, we use techniques described in Section 2.5 to combine multiple frequencies to form a single map with an overall spectral index, as well as the beam-averaged spectral index as a function of time. We pixelize the sky to HEALPIX $n_{\rm side} = 64$. Since the size of the **A**-matrix is proportional to the number of frequency bins, and including the entire data set amkes the size of our **A**-matrix too large, we include only one out of every 5 frequency bins throughout the frequency range of 128.5 MHz to 174.5 MHz. This forms an **A**-matrix of size roughly 6×10^5 by 4×10^4. Since **A** has an order of magnitude more rows than columns, computing $\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A}$ is the speed bottleneck and takes more than 2 days when parallelized on a single CPU. In comparison, inverting $\mathbf{A}^t\mathbf{N}^{-1}\mathbf{A}$ takes about 3 hours.

In order to calculate the relative amplitude at different frequencies, to solve for the re-phasing degeneracies in the data, and to empirically estimate the level of noise and systematics in each data set, we iterate the process described in Section 2.3 until the amplitudes converge to within 0.1% and re-phasings within 0.1°. At each iteration, we calculate visibilities using the solution from the previous iteration as a model, and for each frequency we fit for the re-phasing, the relative amplitude, as well as the overall error RMS. Both the best fit amplitude and the error RMS are used to re-weigh the noise covariance matrices for each frequency. When iterating, we prioritize the map's accuracy in modeling visibilities over its noise properties, so we choose a weaker regularization of $\epsilon^{-1} = 1500$ K. For the final map we choose a stronger regularization with $\epsilon^{-1} = 300$ K to obtain lower noise in the map at a cost of lower resolution in the noisy areas. The map's overall amplitude is locked to the Parkes map at 150 MHz (Landecker & Wielebinski 1970) using the overlapping region. Cyg A, Cas A, and their "ringing" are removed using the CLEAN algorithm (Högbom 1974). The map we obtain together with its angular resolution and error bars are shown in Fig. 4.

4.3 Error Analysis

Fig. 4 shows that the map we obtain agree very well with the prediction of the GSM at 150 MHz. In this section, we discuss the errors in our map and their possible causes in more detail. In terms of overall amplitude, the map's overall amplitude is locked to the Parkes 150 MHz equatorial map (Landecker & Wielebinski 1970), which has 20 K uncertainty in zero level and 4% in temperature scale. In order to compare detailed structures in our result with the Parkes map and the GSM, we calculate the normalized error δ maps, shown in Fig. 4. The median δ compared to the filtered GSM is 2.16, and 2.79 to the unfiltered Parkes map, which are slight higher than what one might expect.

There are a few factors that make the median δ higher than 1. Firstly, our modeling of our instrument is not perfect, leading to error in our **A**-matrix. As we investigated in detail in Zheng et al. (2014), our beam model has up to 10% error in some directions. Our empirically estimated visibility errors are dominated by slowly varying modes, and beam mis-modeling is the most likely cause. Another cause of error is the averaging over frequency, which assumes constant spectral index over the sky. As we show in much more detail in the next section, the spectral index changes by as much as 0.5 from the Galactic plane to out-of-plane regions, so this introduces an error for the edge frequencies on the order of $(175\,\text{MHz}/150\,\text{MHz})^{0.5} - 1 = 8\%$. Moreover, as we have seen in our simulation results in Fig. 1, pixelization can also cause un-modeled error near high temperature regions such as the point sources (Cas A and Cyg A) and the Galactic center. Lastly, neither the GSM nor the Parkes map are the ground truth: the GSM at 150 MHz is essentially an interpolation product between a 45 MHz sky map and a 408 MHz sky map, with an estimated relative error of 10% (de Oliveira-Costa et al. 2008). For the equatorial Parkes map, due to its 2.2° resolution, we cannot apply the PSF matrix to it before comparing it to our result. The lack of PSF on the Parkes map leads to higher error in the comparison, which we think is what makes the median δ higher for the unfiltered Parkes map than for the GSM.

4.4 Spectral Index Results

4.4.1 Spectral Indices from 128 MHz to 175 MHz

In the iterative process to compute the 150 MHz map, we also obtain relative amplitudes between all the data sets at different frequencies. At each frequency, we calculate visibilities using the 150 MHz map we obtained as a model, and compute the relative amplitude by comparing them to our data. We then perform a linear fit in the log(amplitude)-log(frequency) space, and compute an overall spectral index of -2.73 ± 0.11, as shown in Fig. 5. The error bar is calculated using the residuals in the fitting process. In addition to this overall spectral index averaged over the entire data set, we also perform the same procedure on subsets of the data, and fit for spectral indices for every half an hour of LST. The resulting time series of spectral indices varies smoothly between -2.4 and -2.8, as shown in Fig. 5.

In comparison to our spectral index result, the Experiment to Detect the Global EoR Signature (EDGES; Bowman et al. 2008) measured spectral index of -2.5 ± 0.1

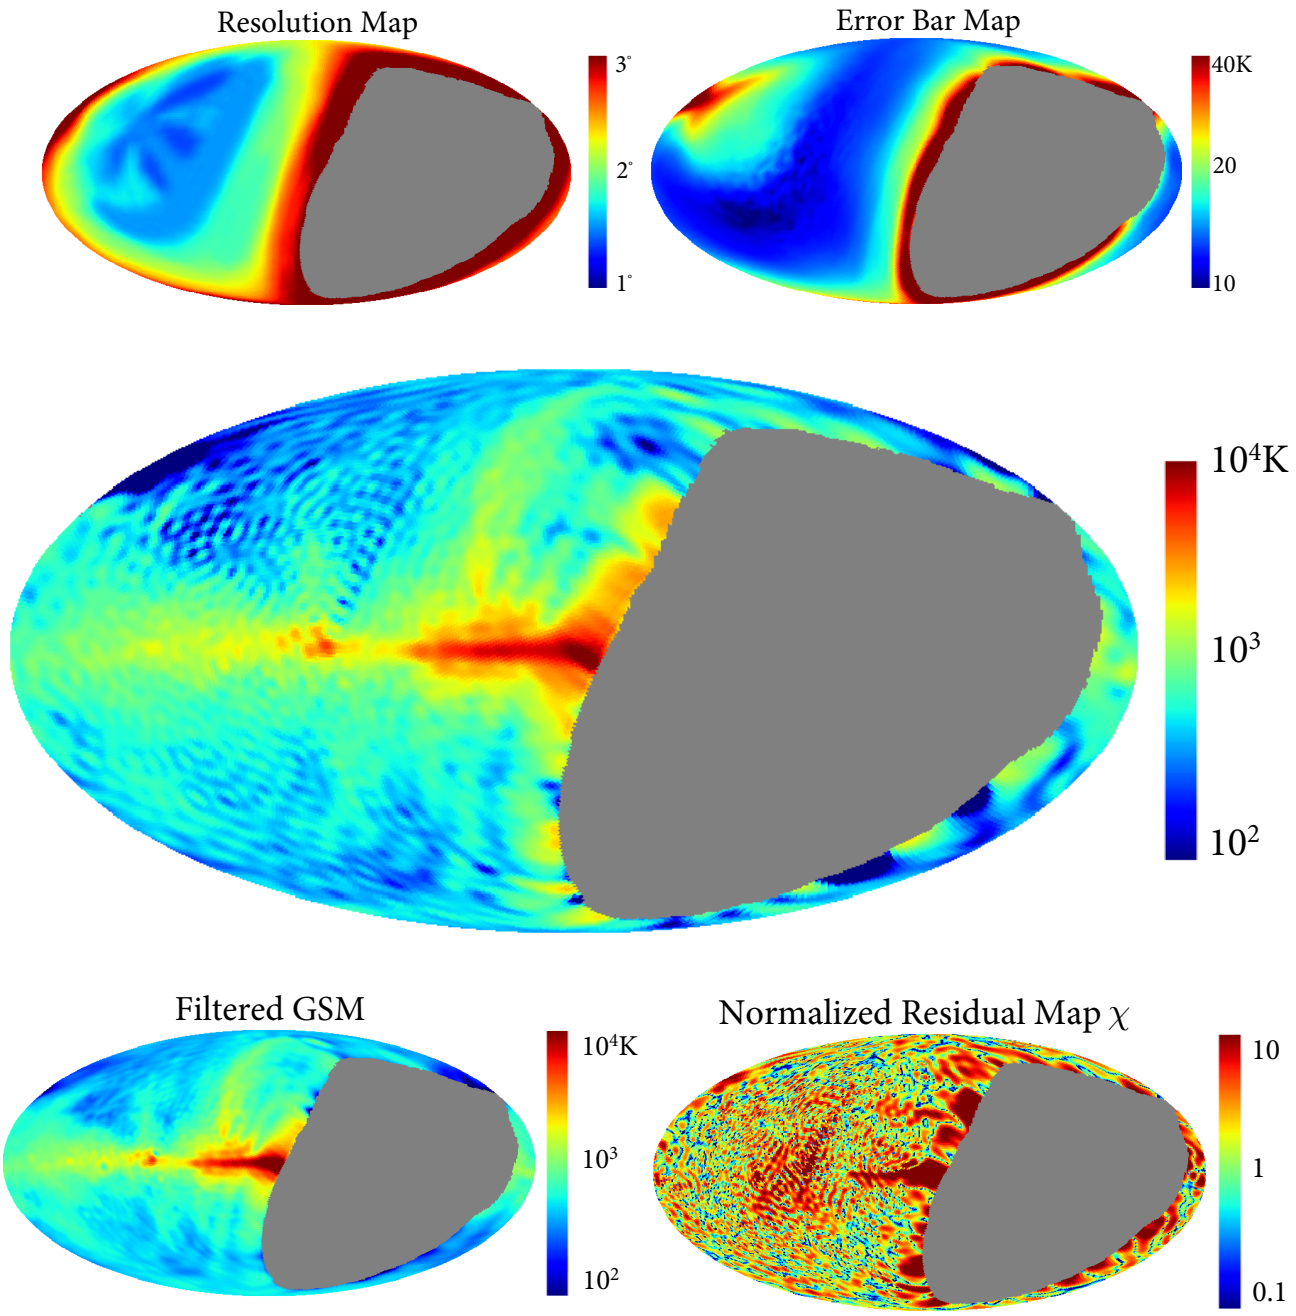

Figure 4. Top left: angular resolution map obtained from FWHM of the point spread functions. Top right: error bar map obtained from $\sqrt{\Sigma_{ii}}$ based on empirically estimated noise covariance $\mathbf{N}$. Mid: the northern sky map at 150 MHz, averaged from 128.5 MHz to 174.5 MHz, with Cas A and Cyg A removed using CLEAN. Bottom left: The GSM at 150 MHz, with the PSF matrix applied. Bottom right: δ as defined in Eq. (13), which represents the ratio between the difference between the maps and our map's error bars. The median δ is 2.16.

(Rogers & Bowman 2008), centering at an out-of-plane region at declination $-26.5°$ and right ascension 2 h. Our overall spectral index agrees with the spectral index obtained by EDGES, with a 1.5σ difference. However, the difference is likely not due to statistical variation alone. The overall spectral index we present is averaged over the northern sky, so we are observing a different patch of sky compared to EDGES. Unlike EDGES, our result is influenced by both Cyg A and Cas A[2]. These strong radio sources have spectral indices of -2.7 and -2.8, respectively (Parker 1968; Vinyaĭkin 2006; Vinyaikin 2014), so they would shift our results towards a steeper spectral index.

[2] The point sources are present in our spectral index results because our spectral index results are obtained in the visibility space, whereas the CLEAN algorithm that removed the point sources is performed in the image space.

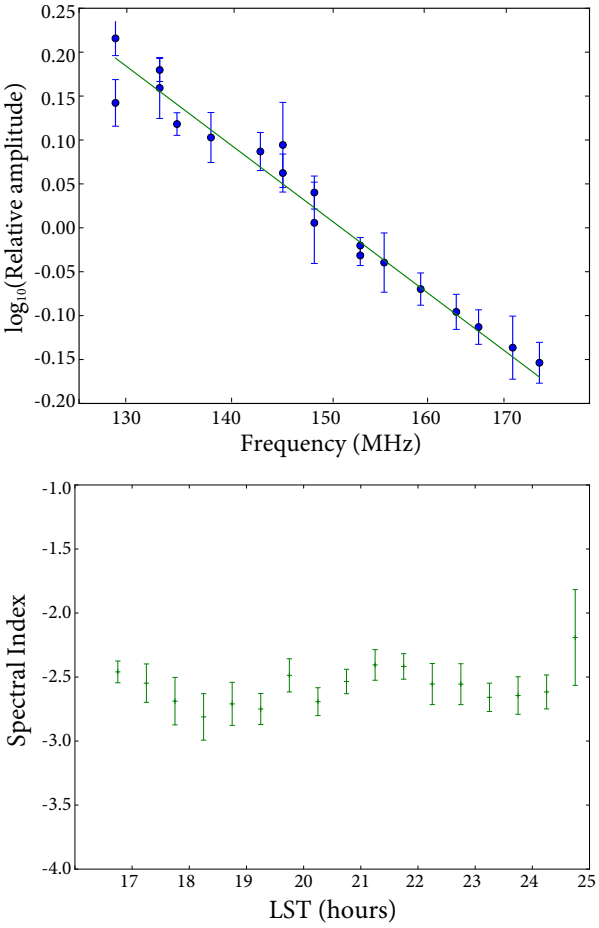

Figure 5. Map-averaged overall spectral index fit (top) and beam-averaged spectral index over LST (bottom). The relative amplitudes and error bars in the top plot are obtained from the iterative procedure described in Section 4.4. For the bottom plot, each LST is observed on 5 different frequencies on 5 different nights (see Fig. 3), so we can obtain a spectral index at each LST by fitting the overall amplitude over frequency with a power law. The error bars are 1σ, empirically estimated using the residuals of the power law fits.

4.4.2 Spectral Indices from 85 MHz to 408 MHz

In addition to spectral indices in the EoR frequency range computed using the MITEoR data alone, we also calculate maps of spectral indices by comparing our MITEoR map to the Parkes map at 85 MHz (Landecker & Wielebinski 1970) and the Haslam map at 408 MHz (Haslam et al. 1981, 1982; Remazeilles et al. 2015). We compute per-pixel spectral index maps for all three pairs of these three maps, as shown in Fig. 6.

The medians of the spectral indices shown in Fig. 6 are $-2.60 \pm 0.29 \pm 0.07$, $-2.43 \pm 0.18 \pm 0.04$, and -2.50 ± 0.07, respectively. The first error bars come from the spread in spectral indices over the sky, and the second error bars come from 4% absolute calibration uncertainty of MITEoR. This is a very weak indication that the spectral index softens over the range from 85 MHz to 408 MHz. For comparison, Platania et al. (2003) presented an overall spectral index of -2.695 ± 0.120 between the 408 MHz (Haslam et al. 1981, 1982; Remazeilles et al. 2015), 1.4 GHz (Reich 1982; Reich & Reich 1986; Reich et al. 2001), and 2.3 GHz (Jonas et al. 1998) maps, which also agrees well with an earlier study at these frequencies in Giardino et al. (2002).

There are two spatial features worth noting in these maps. First, the Galactic plane has softer spectral indices than the out-of-plane regions. The median spectral indices within 5° of the Galactic plane for the three map pairs are -2.27, -2.37, and -2.37, respectively. Softer spectral indices in the Galactic plane are also observed at higher frequencies in Platania et al. (2003), whose spectral index map comes from three maps above the EoR frequency range, as mentioned above.

In addition to softer indices in the Galactic plane, there are regions that clearly deviate from the median near the Galactic poles, such as the blue regions in the Parkes vs MITEoR map, and the red regions in the MITEoR vs Haslam map. Since such departure is not seen in the Parkes vs Haslam map, this suggests that the MITEoR map is about 50 K lower in the Galactic pole regions than what the Parkes and Haslam maps jointly predict. There are two possible causes for this. The first is that the 50 K deficiency in the MITEoR map is due to systematic errors. However, the temperatures in these regions are about 240 K, so neither the 4% absolute calibration uncertainty nor the $\sim$ 15 K error bars can fully explain the 50 K difference. Another possible cause is that the MITEoR map has a higher dynamic range than the Haslam and Parkes maps, so that it recovers more details at the low end of the temperature range compared to these maps. We leave a more careful investigation of this 50 K discrepancy to a future study.

5 SUMMARY AND OUTLOOK

We have presented a new method for mapping diffuse sky emission using interferometric data. We have demonstrated its effectiveness through simulations for both MITEoR and the MWAcore, where we obtained maps with better than 50 K noise and better than 2° resolution for both instruments. We applied this method on the MITEoR data set collected in July 2013, which was absolutely calibrated using Cyg A and Cas A. We obtained a northern sky map averaged from 128 MHz to 175 MHz, with around 2° angular resolution, 5% uncertainty in its overall amplitude, and better than 100 K noise. We also obtained an overall spectral index of -2.69 ± 0.11, and beam-averaged spectral indices that vary over LST between -2.4 and -2.8. Both the MITEoR visibility data and the 150 MHz sky map are available at `space.mit.edu/home/tegmark/omniscope.html`.

As this is our first application of this new method, there are many aspects of it that we are excited to investigate in future work. Throughout this work, we have focused on regularization matrices that are multiples of the identity matrix. However, since the sensitivity varies across the sky, especially in the case of MWAcore, a regularization matrix whose strength varies with sensitivity may achieve a better balance between noise suppression and PSF. It is also interesting to study what the optimal array layout is for imaging diffuse structure, along the lines of Dillon & Parsons (2016). Since the Earth rotates in the East-West direction, we expect the optimal array layout to be very different in the E-W direction than the N-S direction, perhaps similar to that of

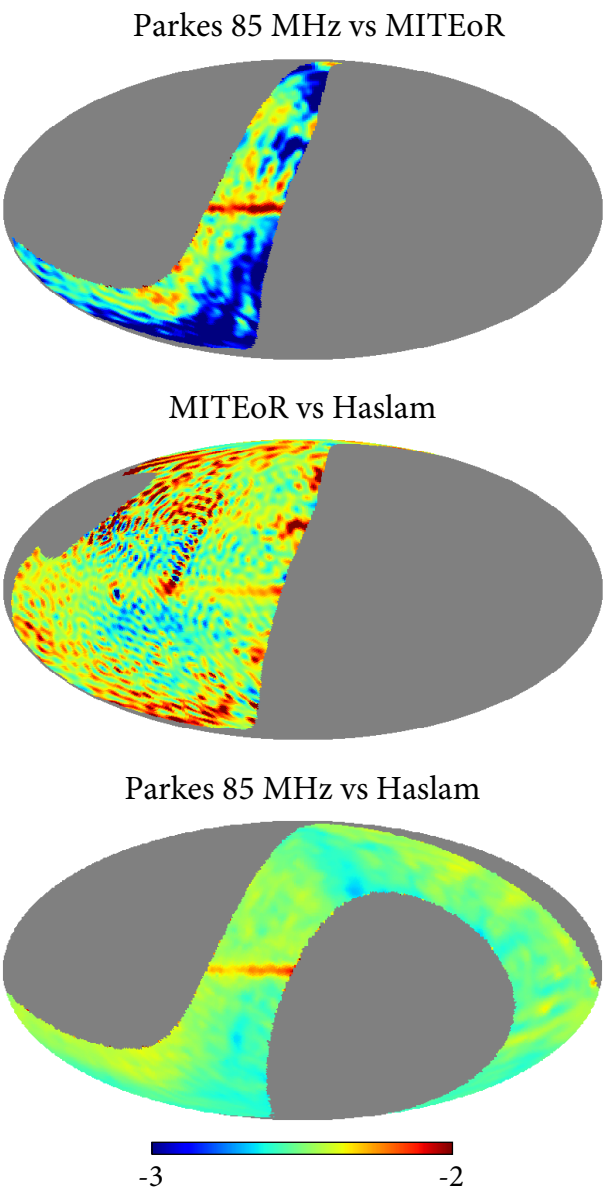

Figure 6. Spectral index maps between the Parkes map at 85 MHz, the MITEoR map at 150 MHz, and the Haslam map at 408 MHz. The MITEoR map is masked for regions with FWHM above 2.5° or error bar above 20 K. The median spectral indices in these maps are −2.60, −2.43, and −2.50, respectively.

PAPER or CHIME (Shaw et al. 2014). Moreover, it is interesting to investigate the effectiveness of this algorithm for instruments with much narrower primary beams, such as HERA. Lastly, it is valuable to perform further study in the effectiveness of this method for the purpose of calibration, such as calibrating the shorter baselines of MWA, which complements the existing calibration algorithms that focus on point source models for very long baselines.

Acknowledgments: MITEoR was supported by NSF grants AST-0908848, AST-1105835, and AST-1440343, the MIT Kavli Instrumentation Fund, the MIT undergraduate research opportunity (UROP) program, FPGA donations from XILINX, and by generous support from Jonathan Rothberg and an anonymous donor. AL acknowledges support from the University of California Office of the President Multicampus Research Programs and Initiatives through award MR-15-328388, and from NASA through Hubble Fellowship grant HST-HF2-51363.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS5-26555. We wish to thank Jacqueline Hewitt and Aaron Parsons for helpful comments and suggestions, Evgenij Vinyajkin for his timely help with our error estimation for the Cyg A spectrum, Dan Werthimer and the CASPER group for developing and sharing their hardware and teaching us how to use it, Richard Armstrong, Matt Dexter, Alessio Magro, Mike Matejek, Qingxuan Pan, Robert Penna, Courtney Peterson, Meng Su, and Chris Williams for help with earlier stage of our hardware development and deployment.

REFERENCES

Ali Z. S., et al., 2015, ApJ, 809, 61
Alvarez H., Aparici J., May J., Olmos F., 1997, A&AS, 124, 315
Barry N., Hazelton B., Sullivan I., Morales M. F., Pober J. C., 2016, preprint, (arXiv:1603.00607)
Bowman J. D., Rogers A. E. E., Hewitt J. N., 2008, ApJ, 676, 1
Dillon J. S., Parsons A. R., 2016, preprint (arXiv:1602.06259)
Dillon J. S., et al., 2014, Phys. Rev. D, 89, 023002
Dillon J. S., et al., 2015a, Phys. Rev. D, 91, 023002
Dillon J. S., et al., 2015b, Phys. Rev. D, 91, 123011
Furlanetto S. R., Oh S. P., Briggs F. H., 2006, Phys. Rep., 433, 181
Ghosh A., Koopmans L. V. E., Chapman E., Jelić V., 2015, MNRAS, 452, 1587
Giardino G., Banday A. J., Górski K. M., Bennett K., Jonas J. L., Tauber J., 2002, A&A, 387, 82
Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, ApJ, 622, 759
Haslam C. G. T., Klein U., Salter C. J., Stoffel H., Wilson W. E., Cleary M. N., Cooke D. J., Thomasson P., 1981, A&A, 100, 209
Haslam C. G. T., Salter C. J., Stoffel H., Wilson W. E., 1982, A&AS, 47, 1
Helmboldt J. F., Kassim N. E., 2009, AJ, 138, 838
Högbom J. A., 1974, A&AS, 15, 417
Jacobs D. C., et al., 2013, ApJ, 776, 108
Jacobs D. C., et al., 2015, ApJ, 801, 51
Jonas J. L., Baart E. E., Nicolson G. D., 1998, MNRAS, 297, 977
Landecker T. L., Wielebinski R., 1970, Australian Journal of Physics Astrophysical Supplement, 16, 1
Liu A., Parsons A. R., Trott C. M., 2014a, Phys. Rev. D, 90, 023018
Liu A., Parsons A. R., Trott C. M., 2014b, Phys. Rev. D, 90, 023019
Maeda K., Alvarez H., Aparici J., May J., Reich P., 1999, A&AS, 140, 145
Mellema G., et al., 2013, Experimental Astronomy, 36, 235
Morales M. F., Wyithe J. S. B., 2010, ARA&A, 48, 127
Offringa A. R., et al., 2014, MNRAS, 444, 606
Paciga G., et al., 2011, MNRAS, 413, 1174
Paciga G., et al., 2013, MNRAS, 433, 639
Parker E. A., 1968, MNRAS, 138, 407
Parsons A. R., et al., 2010, AJ, 139, 1468
Parsons A., Pober J., McQuinn M., Jacobs D., Aguirre J., 2012a, ApJ, 753, 81

Parsons A. R., Pober J. C., Aguirre J. E., Carilli C. L., Jacobs D. C., Moore D. F., 2012b, ApJ, 756, 165
Parsons A. R., et al., 2014, ApJ, 788, 106
Platania P., Burigana C., Maino D., Caserini E., Bersanelli M., Cappellini B., Mennella A., 2003, A&A, 410, 847
Pober J. C., et al., 2013, ApJ, 768, L36
Pober J. C., et al., 2014, ApJ, 782, 66
Pober J. C., et al., 2016, ApJ, 819, 8
Reich W., 1982, A&AS, 48, 219
Reich P., Reich W., 1986, A&AS, 63, 205
Reich P., Testori J. C., Reich W., 2001, A&A, 376, 861
Remazeilles M., Dickinson C., Banday A. J., Bigot-Sazy M.-A., Ghosh T., 2015, MNRAS, 451, 4311
Rogers A. E. E., Bowman J. D., 2008, AJ, 136, 641
Röttgering H., 2003, New Astronomy Reviews, 47, 405
Shaw J. R., Sigurdson K., Pen U.-L., Stebbins A., Sitwell M., 2014, ApJ, 781, 57
Sullivan I. S., et al., 2012, ApJ, 759, 17
Taylor G. B., Carilli C. L., Perley R. A., eds, 1999, Synthesis Imaging in Radio Astronomy II Astronomical Society of the Pacific Conference Series Vol. 180
Tegmark M., 1997, ApJ, 480, L87
Tegmark M., Zaldarriaga M., 2009a, Phys. Rev. D, 79, 083530
Tegmark M., Zaldarriaga M., 2009b, Phys. Rev. D, 79, 083530
Thompson A. R., Moran J. M., Swenson Jr. G. W., 2001, Interferometry and Synthesis in Radio Astronomy, 2nd Edition. New York, Wiley-Interscience, 2001. 692 p.
Tingay S. J., et al., 2013, PASA, 30, 7
Vinyaikin E. N., 2014, Astronomy Reports, 58, 626
Vinyaĭkin E. N., 2006, Astronomy Reports, 50, 143
Wayth R. B., et al., 2015, PASA, 32, e025
Wu X., 2009, in American Astronomical Society Meeting Abstracts #213. p. 226.05
Zheng H., et al., 2014, MNRAS, 445, 1084
de Oliveira-Costa A., Tegmark M., Gaensler B. M., Jonas J., Landecker T. L., Reich P., 2008, MNRAS, 388, 247

APPENDIX A: IMPROVEMENTS TO REDUNDANT CALIBRATION

Zheng et al. (2014) first demonstrated the precision and speed of redundant calibration, which has since been applied to PAPER data analysis in the latest EoR power spectrum upper limits (Ali et al. 2015). In this work we use the same core algorithms as described in Zheng et al. (2014), with improvements that make the algorithm much easier to use. In Appendix A1, we describe the improvement to rough calibration, which is the first step in the redundant calibration process. A similar algorithm has been discussed and used in Parsons et al. (2014). In Appendix A2, we discuss how to reduce the number of undetermined quantities from two per time stamp to two per observing session. The completed redundant calibration software is publicly available at `https://github.com/jeffzhen/omnical`.

A1 Improved Rough Calibration

As described in Zheng et al. (2014), redundant calibration is a three step process: rough calibration, log calibration, and linear calibration. While log calibration and linear calibration do not rely on any sky information, they require rough calibration to get started. Zheng et al. (2014) described a rough calibration algorithm that requires a sky model, which makes the whole pipeline rather cumbersome, especially when dealing with data sets from a new season or a new instrument. In this section we describe a new rough calibration algorithm that can be directly applied to data without any preprocessing, thus making the entire redundant calibration pipeline sky independent. The computational complexity is proportional to the number of antennas.

$\smile_1$	$\longrightarrow$	$\smile_2$	$\longrightarrow$	$\smile_3$	$\longrightarrow$	$\smile_4$
$\downarrow$		$\downarrow$		$\downarrow$		$\downarrow$
$\smile_5$		$\smile_6$		$\smile_7$		$\smile_8$
$\downarrow$		$\downarrow$		$\downarrow$		$\downarrow$
$\smile_9$		$\smile_{10}$		$\smile_{11}$		$\smile_{12}$
$\downarrow$		$\downarrow$		$\downarrow$		$\downarrow$
$\smile_{13}$		$\smile_{14}$		$\smile_{15}$		$\smile_{16}$

Table A1. A 4 by 4 antenna array on a regular grid. Each $\smile_n$ represents an antenna, and each arrow represents a baseline used in the rough calibration algorithm.

Here we describe our algorithm using an 4 by 4 array on a rectangular grid as shown in Table A1, but the algorithm is easily generalized to other redundant configurations. At a given time and frequency, we start with $16 \times 15/2 = 120$ visibilities, v_{ij}, whose phases are γ_{ij}, where i, j are antenna numbers. The goal of rough calibration is to obtain antenna calibration phases, ϕ_i, for all 16 antennas. The simple equation that describes phase calibration is

$$\gamma_{ij} = -\phi_i + \phi_j + \theta_{i-j} \tag{A1}$$

where θ_{i-j} is the true phase shared by all redundant baselines that share the same baseline type with γ_{ij}, and we omit 2π wrapping for this section.

We start by taking advantage of the three degree of freedom in phase degeneracies that are intrinsic to redundant calibration, namely an overall phase to all antenna phases, and two phases corresponding to rephasing the array. Due to these degeneracies, we are free to declare that $\phi_1 = 0$, $\theta_{1-2} = \gamma_{1-2}$, and $\theta_{1-5} = \gamma_{1-5}$ [3]. This then gives us

$$\phi_1 = \phi_2 = \phi_5 = 0. \tag{A2}$$

With the first five phases $\phi_1, \phi_2, \phi_5, \theta_{1-2}, \theta_{1-5}$ all determined, we can now proceed to solve all 16 antenna phases. By applying Eq. (A1) on baseline $\gamma_{2,3}$, we can solve for ϕ_3 since ϕ_2 and $\theta_{2-3} = \theta_{1-2}$ are known. We can repeat the process to obtain ϕ_4. Similarly, we can obtain ϕ_9 using $\gamma_{5,9}$, and extend that to obtain ϕ_{13}. Now we see that after determining the first five phases, solving for all the antennas is simply a matter of traversing the entire array and visiting every antenna with either one of the two baseline types we picked, in this case θ_{1-2} and θ_{1-5}.

As shown in Table A1, the only baselines we used in this example are the arrows in the table, so the computational complexity is $\mathcal{O}(N)$ where N is the number of antennas. If noise is a concern, one can increase the number of baseline types used to visit each antenna, so each phase calibration is sampled multiple times using multiple baseline types. For both MITEoR and PAPER 32 element data (Parsons et al. 2014), we found that using just the two most redundant baseline types suffices.

It's important to note that while this rough calibration

[3] We choose θ_{1-2} and θ_{1-5} because they are the two shortest non-parallel baseline types.

is much faster than log calibration and linear calibration, it is not advisable to re-run it on every data set before log calibration. As we pointed out above, rough calibration makes an arbitrary decision on the phase degeneracies, which later requires absolute calibration (with an accurate sky model) to determine. If each time step (typically a few seconds) uses a different rough calibration, we will have to run absolute calibration separately on every time step, which is very challenging as we may not always have good calibrators in the sky. It is much easier to run rough calibration only once at a single time stamp, and use it for a long period of data, such as one night or even a whole season. This way the entire time window will share one constant set of phase degeneracies, which are easy to determine in absolute calibration. Similarly, log calibration and linear calibration also make arbitrary choices of these same re-phasing degrees of freedom, so it's important to project out any re-phasing components from the calibration phase solutions, which we discuss in the next section.

A2 Removal of Temporal Fluctuation in Phase Calibration Degeneracies

Both log calibration and linear calibration have the same set of phase degeneracies that cannot be determined by the self-consistency of the redundant measurements in a redundant array. Without careful treatment, these degeneracies can vary over time and become very difficult to determine. In this section we describe how we project out the temporal fluctuation in phase degeneracies in our redundant calibration solutions.

We start by reiterating the mathematical framework we first described in Zheng et al. (2014). When calibrating visibilities at any given time stamp and frequency bin, the quantity that log calibration minimizes is

$$\sum_{jk} |(\theta_{j-k} - \phi_j + \phi_k) - \gamma_{jk}|^2, \tag{A3}$$

whereas linear calibration minimizes

$$\begin{aligned}&\sum_{jk} |(y_{j-k} g_j^* g_k) - v_{jk}|^2 \\ =&\sum_{jk} \left||y_{j-k} g_j^* g_k| \exp\left[i(\theta_{j-k} - \phi_j + \phi_k)\right] - v_{jk}\right|^2 .\end{aligned} \tag{A4}$$

For any given set of θ_{j-k} and ϕ_k, one can add an arbitrary linear field $\boldsymbol{\Phi}\cdot\boldsymbol{r}_j + \psi$ to ϕ_j across the entire array, and simultaneously subtract $\boldsymbol{\Phi}\cdot\boldsymbol{d}_j$ from θ_{j-k}, without changing the minimized quantities:

$$\begin{aligned}\theta'_{j-k} - \phi'_j + \phi'_k \equiv&(\theta_{j-k} - \boldsymbol{\Phi}\cdot\boldsymbol{d}_{j-k}) - (\phi_j + \boldsymbol{\Phi}\cdot\boldsymbol{r}_j + \psi) \\ &+ (\phi_k + \boldsymbol{\Phi}\cdot\boldsymbol{r}_k + \psi) \\ =&\theta_{j-k} - \phi_j + \phi_k.\end{aligned} \tag{A5}$$

Here $\boldsymbol{d}_{j-k} \equiv \boldsymbol{r}_k - \boldsymbol{r}_j$ is the baseline vector for the unique baseline with best-fit visibility y_{j-k}. Thus, the quantities in expressions A3 and A4 that the calibrations minimize are degenerate with changes to the linear phase field $\boldsymbol{\Phi}$ and the scalar ψ. This gives us 4 degenerate phase parameters: one overall phase ψ (which has no impact on visibilities) and three related to the three degrees of freedom of the linear function $\boldsymbol{\Phi}$ (which reduces to two for a planar array).

From the discussion above, we have two phase degeneracies per set of redundant visibilities at each time stamp that must be determined using other methods. Solving for these degeneracies at every time stamp can be very challenging, if not infeasible, given the large field of view of the instruments and limited availability of high quality calibrators. Therefore, we need to project out the temporal variability of the degeneracies so that we only have two degeneracies per observing session. To do this, we first pick a time stamp t_0 near the middle of the observing session, where the phase solutions for all antennas form a vector $\boldsymbol{\phi}_0$. Then, for any other time t_1 and solution $\boldsymbol{\phi}_1$, we project out the degeneracy difference between them by calculating

$$\boldsymbol{\phi}'_1 = \mathbf{P}\boldsymbol{\phi}_0 + (\mathbf{I} - \mathbf{P})\boldsymbol{\phi}_1, \tag{A6}$$

where the projection matrix $\mathbf{P}$ into the degeneracy space is defined as

$$\mathbf{P} \equiv \mathbf{R}(\mathbf{R}^t\mathbf{R})^+\mathbf{R}^t, \tag{A7}$$

and $\mathbf{R}$ is an $n_{\rm ant}$ by 4 matrix whose ith row is $[\boldsymbol{r}_i\, 1]$, $\sum_i \boldsymbol{r}_i = \mathbf{0}$[4], and $^+$ denotes pseudo-inverse. In this set-up, the phase degeneracy space is the column space of $\mathbf{R}$, and the degeneracy coefficients $(\boldsymbol{\Phi}, \psi)$ are the coordinates in this degeneracy space. Regardless of the choice of t_1, $\boldsymbol{\phi}'_1$ always has the same component in the degeneracy space, $\mathbf{P}\boldsymbol{\phi}_0$. We thus replace $\boldsymbol{\phi}_1$ by $\boldsymbol{\phi}'_1$ as our new phase calibration solution.

While using $\mathbf{P}$ to project out degeneracies removes temporal variability in the degeneracies, it has the drawback of projecting out the 'degeneracy component' in the variation of true phase solutions at the same time. Fortunately, since the projection is from a $n_{\rm ant}$-dimensional space to a 2-dimensional space, the error caused by projection is much smaller than the phase drift in the calibration solutions.

To be more quantitative, let $\boldsymbol{\phi}_{\rm var}$ be the difference vector in true phase calibration solutions between t_0 and t_1, $\boldsymbol{\phi}_{\rm err}$ the part of $\boldsymbol{\phi}_{\rm var}$ that is erroneously projected out, $\mathbf{R}_0$ an $n_{\rm ant}$ by 4 matrix whose ith row is $[\boldsymbol{r}_i\, 0]$, and $\mathbf{P}_0 \equiv \mathbf{R}_0(\mathbf{R}^t\mathbf{R})^+\mathbf{R}^t$[5]. We then have

$$\boldsymbol{\phi}_{\rm err} = \mathbf{P}_0\boldsymbol{\phi}_{\rm var}, \tag{A8}$$

and

$$\langle\boldsymbol{\phi}_{\rm err}\boldsymbol{\phi}^t_{\rm err}\rangle = \mathbf{P}_0\langle\boldsymbol{\phi}_{\rm var}\boldsymbol{\phi}^t_{\rm var}\rangle\mathbf{P}^t_0. \tag{A9}$$

Assuming the drift for each antenna is uncorrelated and has a typical magnitude of ϵ, the typical magnitude of phase projection error can be described by the median of the diagonal entries of $\epsilon\sqrt{\mathbf{P}_0\mathbf{P}^t_0}$. For an 8 by 8 regular grid array such as MITEoR, median of the diagonal of $\sqrt{\mathbf{P}_0\mathbf{P}^t_0}$ is 17.6%, and for a 16 by 16 array this becomes 8.8%, for a 32 by 32 array 4.3%. This result is unsurprising, since we would expect the error fraction to be roughly $\sqrt{2/n_{\rm ant}}$ using the dimensionality argument alone. Therefore, the error caused by the

[4] Picking the origin for $\boldsymbol{r}_i$ to be center of the array simplifies the calculation in the next paragraph, with no impact on any visibility results.

[5] When $(\mathbf{R}^t\mathbf{R})^+\mathbf{R}^t$ is applied to the calibration solutions, it gives the degeneracy coefficients $(\boldsymbol{\Phi}, \psi)$. We can use these coefficients to convert back to the actual degenerate components of the solution, but we use $\mathbf{R}_0$ rather than $\mathbf{R}$ because we do not want to penalize ourselves for the overall degeneracy ψ that does not impact visibilities.

projection operation diminishes as the number of antennas grows.

Lastly, an important subtlety worth noting is phase wrapping on $(-\pi, \pi]$. All the equations in this section are dealing with phase, so they are subject to phase wrapping. Phase wrapping can potentially render Eq. (A6) invalid. Therefore, rather than calculating Eq. (A6) directly, it is much less likely to have phase wrapping by calculating

$$\boldsymbol{\phi}_1' = \boldsymbol{\phi}_0 + (\mathbf{I} - \mathbf{P})(\boldsymbol{\phi}_1 - \boldsymbol{\phi}_0). \tag{A10}$$

Unlike $\boldsymbol{\phi}_0$ and $\boldsymbol{\phi}_1$ which can have any values from $-\pi$ to π, $\boldsymbol{\phi}_1 - \boldsymbol{\phi}_0$ should be small angles around 0. For large arrays, even $\boldsymbol{\phi}_1 - \boldsymbol{\phi}_0$ will have large angles that wrap around π for antennas near the outer part of the array. In this case, one can first take only the inner part of the array to find a rough solution to the degeneracy coefficients, and remove them from them from $\boldsymbol{\phi}_1 - \boldsymbol{\phi}_0$ to reduce the angles therein and eliminate phase wrapping, and then carry out Eq. (A10).

APPENDIX B: DYNAMIC PIXELIZATION

Multiplications and inversions of 10^5 by 10^5 matrices can currently be done in a matter of days, but since the computational times grows as $\mathcal{O}(\Delta\Omega^{-3})$, improving the resolution can be very computationally demanding. To reduce the matrix size. we can resort to a more intelligent pixelization scheme, which we call dynamic pixelization.

Having uniform pixelization can be wasteful. For example, near the edge of our observable region, fine pixelization brings little advantage since our sensitivity is very low. In contrast, regions with strong flux demand much finer pixelization, such as the Galactic center and strong point sources such as Cas A and Cyg A. We use a recursive algorithm to generate a pixelization scheme that accommodates all of the above considerations. We start with a map predicted by the GSM[6], and a uniform but coarse HEALPIX nested pixelization with $n_{\rm side} \sim n_\lambda$ for the longest baseline[7]. We then calculate a beam weighted sky map $s_{\rm bweight}$ given by

$$s_{\rm bweight}(\hat{\boldsymbol{k}}) = \int_t s(\hat{\boldsymbol{k}}) B(\hat{\boldsymbol{k}}, t) dt. \tag{B1}$$

Lastly, we recursively divide each pixel into 4 sub-pixels according to HEALPIX nest scheme, until the standard deviation among the $s_{\rm bweight}$ values at these 4 sub-pixels are below a threshold value σ. Intuitively, since each pixel corresponds to a column in the $\mathbf{A}$-matrix, the algorithm only splits each column into 4 when the resulting 4 columns will be sufficiently different in fitting the data.

We decide the threshold σ by numerically searching for the optimal choice. As larger σ translates into fewer pixels but less pixel precision, we choose the largest σ such that the simulated pixelization error is less than 1% in power compared to the noise in the visibility data. Fig. B1 shows an example pixelization scheme for MITEoR data at 160 MHz, where the final pixel count is less than a quarter of uniform pixelization, leading to a factor of 64 speed-up in computation time, while the pixelization error remains negligible compared to noise in the data.

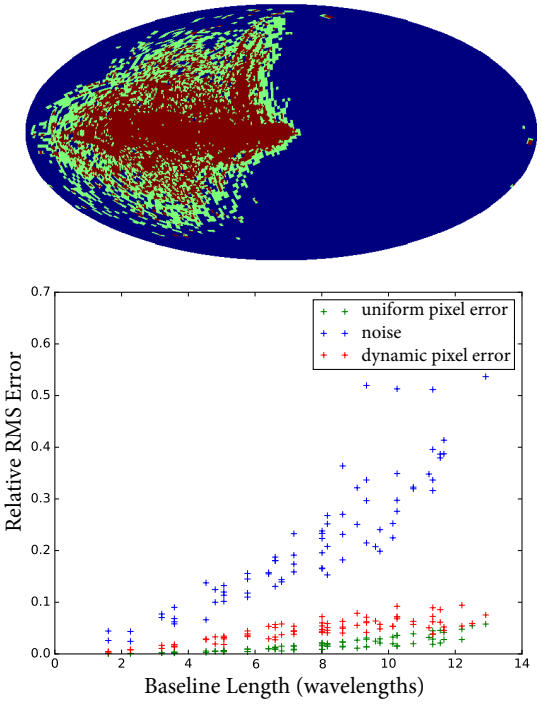

Figure B1. The upper plot shows a dynamic pixelization scheme, where the blue, green, and red regions represent $n_{\rm side}$ of 32, 64, and 128, respectively. This scheme contains 48900 total pixels, 24.9% of the pixel count in a uniform $n_{\rm side} = 128$ pixelization. The lower plot shows the errors in the simulated visibilities compared to those simulated with a $n_{\rm side} = 512$ pixelization, and we see that the dynamic pixelization error is a factor of two higher than uniform pixelization, and it is negligible compared to the noise level assuming one night's observation.

[6] Note that using the GSM for pixelization does not automatically mix GSM information into our solution.

[7] This relation is derived from $\frac{4\pi}{n_{\rm pix}} \sim \left(\frac{1}{n_\lambda}\right)^2$.